

*Amendments to the Claims*

The listing of claims will replace all prior versions, and listings of claims in the application.

Claims 1-7. (Cancelled).

Claim 8. (Currently amended) In an information retrieval system having a current query, and a search engine that is responsive to a query to retrieve documents generally relevant to the query from a document collection, a method for refining the current query, the method comprising:

receiving information regarding the relevancy of documents retrieved at least in part in response to the current query;

ranking the retrieved documents at least in part in accordance with the relevancy information;

forming at least one candidate query based at least in part on one or more of:

the rankings and,

analysis of locations of the retrieved documents in [[the]] a latent semantic index vector space formed from at least the retrieved documents;

applying at least one candidate query to the document collection;

ranking the documents retrieved in response to each applied candidate query in accordance with the received relevancy information;

comparing the ranking of documents retrieved in response to at least one applied candidate query and the ranking of documents retrieved in response to the current query with the received relevancy information; and

choosing the query, from among the current query and each candidate query, which produces the best ranking.

Claims 9-11. (Cancelled).

Claim 12. (Currently amended) A method for refining a query, the method comprising:

receiving information regarding the relevancy to an information need of documents retrieved from a document collection at least in part in response to a current query;

ranking the retrieved documents at least in part in accordance with the relevancy information;

forming at least one candidate query based at least in part on at least one of:

the rankings and,

the locations of the retrieved documents in [[the]] a latent semantic index vector space formed from at least the retrieved documents;

applying the at least one candidate query to the document collection;

ranking the documents retrieved in response to each applied candidate query at least in part in accordance with the received relevancy information;

comparing the ranking of documents retrieved in response to at least one applied candidate query and the ranking of documents retrieved in response to the current query with the received relevancy information; and

choosing the query, from among the compared queries, which produces the best ranking.